SYSTEM FAULT PROTECTI ON DES I GN FOR THE CASS IN I SPACECRAFT

John P. Slonski
Jet Propul sion Laboratory
California Insti tute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109
818-3 54 -2844

Abstract- -Fault protection can i nclude a wide range of topi cs, rangi ng from faul t preventi on to autonomous faul t detection and recovery. 'l'hi s paper wi 11 address a porti on of the autonomous f aul t det ect ion and recovery impl emented onboard the Cassi ni spacecraft . Specif ical ly, t-he topic i s system] evel fault protection design, as opposed to subsystem f aul t protecti on desi gn.

The design of system fault protection f or the Cassi ni spacecraft wil l be described at. a hi gh 1 evel in order t o def i ne the guiding principles of the design. Thi s wi ll include exami ning driving requi rements, high level design trades, and major archi tectural e] ements, incl uding practical details of their design. Finally, a detai led design description will be given of two Cassini fault protec ti on responses whi ch are] ikel y t o be used on most other spacecraft .

1. 1 NTRODUCTI ON

Mi ssi on and Spacecraft Overvi ew

The Cassini mission will explore the Saturni an system in general , and i ts moon Ti tan in parti cul ar, beginning after spacecraft arrival in 2004 . A compl ement of science equipment located on both a Ti tan at. mospheri c entry probe and the pri mary space craf t wi 11 accompl i sh thi s over the course of a four year orbital tour. The mission is a joi nt. proj ect undert aken by NASA and the European Space Agency (ESA) . The entry probe and i ts support avionics are provided by ESA, and the. remainder of t he space craf t i s provi ded by NASA . Sci ence equi pment on each is provi ded from both U. S. and European sources . Instruments specific to studyi ng i nf rared, visible, and ul traviol et. spectral ranges, as wel] as various field, energeti c particle, plasma, dust, and wave phenomena are bei ng flown along with instruments capable of sampl ing atmospheric const i tuent s and acquiri ng visi bl e and RADAR i mages . Toget her, the scientifi c investigate ons wi 11. combi ne to focus on Saturn' s ring and sate]] ite system, atmosphere, magnetosphere, as well as Ti tan' s cl ouded atmosphere and surface .

The spacecraft, including instruments and probe, wi ll weigh 5600 kil ograms (12, 400 pounds) when] aunched by a Ti tan/Centaur i n October of 1 997 . Because of i ts large mass, the spacecraft must acqui re addi tional energy by gravi tati onal assi sts from Venus, Earth, and Jupi te r duri ng a seven year transi t L o Saturn. The spacecraft physi cal 1 y consi sts of a 4 meter diameter high gain antenna, a 12 si ded ring-like bus structure housing most of the

engineering and science electronics, a propulsion module carrying 3100 kg of propellants, the Titan probe, three radioisotope thermoelectric generators (RTGs), four attitude control thruster clusters, four reaction wheels, two main rocket engines, an 11 meter magnetometer boom, three 10 meter radio and plasma wave science antennas, and a dozen other science sensors mounted in various locations.

Engineering subsystems provide for telecommunication links operating at X-Band, the generation and distribution of RTG power, command distribution and data handling employing two redundant 1750A processors, a 1553B data bus, and twenty-six engineering and science remote terminals, bulk data storage on two 1.8 gigabit solid state recorders, and attitude control built around two redundant 1750A processors and various sensors and actuators. Additional functions include thermal control using both electrical and radioisotope heaters, and mechanical actuations by means of motors and pyrotechnic devices.

Fault Protection

Equipment failures occur on spacecraft (S/C), in spite of the best efforts to execute reliable design, use high reliability parts, and thoroughly test everything at the assembly, subsystem, and system levels before committing to launch. Many of these failures are of a non-threatening nature and/or can be eventually worked around by ground intervention or changes in mission plans. Many other failures, however, place the S/C or the mission in immediate jeopardy, or are very difficult for ground personnel to diagnose and correct. Add to this the very long Cassini mission (at least 11 years), and there is a critical need to design and implement on board fault detection and recovery processes, and to do so in a thorough, effective, and practical. way. The first step is to understand the requirements.

2. DRIVING REQUIREMENTS

The highest-driving requirement which Cassini System Fault Protection (SFP) responds to is the Project's single point failure policy: "No credible single point failure shall prevent attainment of the objectives listed below, or result in a significantly degraded mission, . . .
a) . . . minimum essential engineering data and command capability. . .
b) Successful. Earth swingby
c) Successful . . . Saturn Orbit Insertion
d) Successful. Huygens Probe delivery and data return
e) . . . science data from all except one instrument . . .11

Exemptions were granted for items whose failure is not credible due to the existence of large design margins, such as primary structure, propulsion tanks and lines, and RTG power sources. A listing of all singlepointfailure policy exemptions is shown in Table 1.

Although the highest driving requirement on SFP is to provide protection for single faults, there is a goal to recover from multiple faults, provided the faults are located in independent fault containment regions. Obviously, two independent faults occurring at widely separated times in the mission are easily handled. However, two simultaneous faults, or a second fault occurring before the response to t-he first has completed, cannot. always be protected against. As discussed below, the approach to handling these possibilities was one of the key issues addressed in defining the design architecture.

Continuing with the driving requirements, onboard responses must be designed wit-h the following priorities:
1) Protect critical S/C functionality
2) Protect. S/C performance and consumables
3) Minimize disruptions to normal sequence operations
4) Simplify ground recovery response, including providing for downlink telemetry

Implicit in 1), protection of critical S/C functionality, is providing for uplink commendability, import-ant. not only for the conduct of the mission, but also to provide for ground response to unforseen anomalies or complications that can only be covered after realtime evaluation. For this reason, all responses ensure reliable uplink communications performance following an anomaly. There is also an specific response which will restore command-ability if no valid commands have been received within a specified period of time.

Additional driving requirements include:
a) Maintain a safe state following an anomaly for two weeks
b) Provide fault detection, ion with only engineering data, i.e. be independent of science instrument. available
c) Ensure autonomous completion of mission critical sequences
d) Provide diagnostic data after a fault response sufficient to reconstruct the sequence of fault protection actions
e) Provide for protection of any recorded data leading up to the anomaly and the recording of data following the anomaly response
f) Accommodate ground testing without risk of S/C damage

3. ALLOCATION OF FAULT RECOVERY RESPONSIBILITY

As shown in Figure 1, fault recovery responsibility is divided in the first place between the ground team and the S/C. A fundamental constraint in meeting the single point failure policy is that ground response cannot be assumed for the most critical failures, primarily because of' large Earth-S/C distances, and also because the mission and S/C designs impose long communication outage periods. Therefore, time-critical responses to faults must. be provided autonomously onboard the S/C. In general, autonomous fault protection is included onboard the S/C if a ground response is not. feasible or practical, or if action is required within two weeks of detecting the failure.

Responsibility for autonomous S/C fault recovery is further divided among subsystems and SFP. Subsyst ems are in general responsible for providing recovery of their own functionality, if they have the capability to do so. If not, or if there is a part of the recovery which requires a non-standard operation by another subsystem, SFP becomes involved.

Not-e that autonomous S/C fault protection does not include measures used in the handling of data to protect. against bit errors. Data error detections may be used by autonomous S/C fault protection, however, to detect and correct fault conditions.

4. SYSTEM FAULT PROTECTION FUNCTIONS

The functions which SFP monitors and protects are in t-he general areas of radio frequency (RF) communication, power, thermal state, onboard computer viability, and propellant pressurization. The RF communications functions are: uplink commandability, RF exciter output, and RF power amplifier (travel ing wave tube) output, . SFP has a support role with respect to recovery from power faults. Primary responsibility in the event of power demand exceeding capabi lity is allocated to subsystem hardware, which turns off non-essential loads, with SFP following up by establishing a safe S/C state . Thermal fault protection is restricted to correcting emergency conditions brought on by excessive solar illumination of apertures and radiators . Onboard computer viability is monitored by SFP checking for the regular generation of "heartbeats". Propel lant tank pressures are monitored in order to protect against. a gross pressure regulator failure. In addition, SFP responds to specific requests to support the attitude control subsystem (AACS) , including providing additional power margin for its internal fault protection actions.

5. KEY ARCHITECTURAL DESIGN ISSUES

There were three issues which were key to defining the SFP architecture. The first involved the response of SFP in the event- of multiple simultaneous faults. Although the Project's single point failure policy does not require SFP to protect against more than one fault, there is a possibility that during the response to a fault, a second independent fault could occur, especially if the initial fault response is a lengthy one. It. is a stated goal to recover from such double failure scenarios. In addition, a single fault can result in multiple effects, with each symptom looking like an independent fault requiring a different response . It is a hard requirement to recover from this scenario. On Cassini, the alternatives for handling multiple simultaneous faults or fault symptoms were to 1) work one fault. at a time in a first come/first served manner, 2) work one fault at a time with a priority structure allowing more important fault. responses to interrupt less important. ones, or 3)' let different faults be responded to simultaneously. Previous SFP designs nave used the structurally simple third approach, and deal t wi th i nteract i on problems ad hoc, e.g. having a response temporarily disable another potentially

interfering response. Although previous experience wit-h the third approach did not exhibit any major problems, there was concern over having to find and deal with potentially complex response interactions. There was also the problem of validating non-interference of any two fault responses running simultaneously. The first approach has the flaw of delaying the execution of a time critical fault response until the completion of a lengthy non-critical response which was initiated sooner. This left the priority driven one-at--a-time architecture, which is more complex to implement-, but avoids the pitfalls of the other two options.

In order to simplify testing, it was desirable to minimize the number of priority levels. Instead of assigning a unique priority to each fault-, more than one fault. would share the same priority. Faults sharing the same priority level would be worked on a first-come basis. After several iterations of assigning faults to priority levels, it was determine] that this scheme would work with four separate priority levels. Thus , the final design ended up being a combination of options 1 and 3.

The second issue involved simultaneity, also. In this case, the issue was whet-her or not onboard stored sequences, i.e. scripts, should continue running during and after the execution of a fault response . Stored sequences and fault responses can potentially interfere with each other. To avoid any chance of this, one alternative was that stored sequences could be cancelled prior to the first action of any system fault response. (Only a sequence marked "critical " would be restarted following the response completion .) On t-he other hand, it was desirable to minimize disruptions to normal sequence operations (see response priority 2 above) . Therefore, the selected approach was that fault responses would be pre-evaluated for their compatibility with expected flight sequences, and non- interfering fault responses would be allowed to execute in parallel with them. Interfering fault responses would first stop the sequences so as to prevent interference to or from the remainder of the response. This approach was designed to allow the categorization of a given non- interfering response to be changed inflight, if actual flight sequences changed the potential for interference. (l'he concept of allowing sequences to continue during non-interfering fault responses was later dropped, however, in order to simplify and descope the software design effort. The Cassini design is now that stored sequences are halted prior to any system fault response execution, ant] only critical sequences are restarted after response completion. It is felt that fault response occurrences will not be so frequent- that restarting non-critical sequences will be overly burdensome to the flight- Learn.)

'l'he third issue, how to handle sequence dependencies, was not actually fully explored during the initial phases of the design. It was, however, considered at various times after the design was established. This issue is concerned with how to handle fault responses, or parts of fault responses, which are not appropriate during the execution of certain sequences. As an example, the lack of RF output does not indicate a fault when the RF amplifier has

been turned off by the sequence. Another example is that the safing response should not alter the S/C attitude just prior to or during the Saturn orbit insert-ion burn. An ideal design takes cues from the overall S/C state and the sequence and combines them logically to do the right- thing. The alternative is to require the sequence designers themselves to include special commands to disable inappropriate fault responses and set flags which select desired branches within responses. This latter method was the initial. direction taken by the design. It has not changed, because the ideal design requires too much sophistication to always reliably do the right thing. For example, taking cues from a commanded state register is unreliable when that. register itself suffers a failure. A comprehensive S/C configuration manager, tied in wit-h the fault protection system, could possibly have been designed to work in a reliable, fail-safe way, but that would have been outside the scope originally funded for system fault protection. Enables, disables, and branch flags have instead been defined, along with rules for sequencers to follow. This still appears to be t-he simpler and faster way to go, as long as the rules are simple, and there are not too many of them.

6. OVERALL DESIGN ARCHITECTURE

The Command and Data Subsystem (CDS) is the host. for SFP. CDS functionality, and therefore SFP availability, is ensured by internal CDS fault protection, which has available to it a redundant CDS which can assume the prime role. If the prime CDS fails to satisfy internal. tests, it will reset, and the backup will become prime. However, because the backup CDS is powered off during a majority of the mission, there is a watchdog timer in the Power and Pyro subsystem (PPS) which will turn on the backup CDS if the prime CDS fails to reset the timer at least once every 32 seconds .

The PPS watchdog is part. of a system of watchdogs which ensures the viability of all the engineering subsystem processors on the S/C. Figure 2 shows that chain, which consists of t-he F'PS hardware monitoring for the presence of a CDS watchdoq reset. signal, CDS monitoring of the AACS processors via a heartbeat. signal , and CDS monitoring of the Radio Frequency Subsystem (RF'S) processor via another heartbeat.

Within the CDS software, the SFP consists of a structure Of fault tnonitors, responses, and a manager for coordination. A system failure will result in the related monitor evaluating the fault data, which will then lead to the monitor calling for execution of the appropriate response or responses. Any active response may itself activate a supporting response, within the constraint that no more than eight responses may be act ive at any given time. 'l'he Fault Protection (FP) manager provides for various functions involving the initialization of monitors and responses, the activation of responses, and the resolution of priority conflicts during multiple fault scenarios. Figure 3 contains a list of all t-he SFP monitors and related responses.

The monitors are algorithms coded in ADA, and the responses are algorithms coded in a sequencing language, the same as the one used for stored mission sequences. Monitors basically compare key data (FP data) collected by the CDS from S/C subsystems against defined failure thresholds. Where possible, the data is initially screened to reject anything coming from a failed sensor ("reasonableness" checking). If a data sample passes this screen and is found to exceed the failure threshold, a fault occurrence counter is incremented. Consecutive increments of this counter will eventually result in the counter threshold being exceeded ("persistence checking"), and the monitor will request a response.

After activation by the FP manager, responses basically function in the same way as mission sequences, issuing a pre-defined sequence of commands. One significant difference characterizing fault responses is the extensive use of logical operations, which are used to implement the branching discussed above. Once activated, response sequences are designed to execute to completion without further inputs from the calling monitor. If fault recovery requires trying different actions until successful, the response algorithm select-s a new action for each succeeding entry. After an unsuccessful action, the fault must be redetected by the monitor, which results in the response being requested again.

The FP Manager causes monitors to be executed once per second. Whenever a monitor requests a response, the Manager will cause stored mission sequences to be stopped and the requested response to be activated. In the event that two or more monitors request responses in the same one second management- cycle, or if a monitor requests a response while a previously requested response is still active, the Manager will ensure the highest priority fault takes precedence. The mechanics of how the Manager does this will be discussed later in more detail.

An additional task for the Manager is to initialize monitors and responses. Monitor initialization consists of clearing new or outstanding response requests and resetting internal persistence counters, which occurs as follows. Whenever a monitor or its corresponding response is enabled, the monitor is initialized. (It is important to clear the monitor when a response is enabled, so that conditions which existed only while it was disabled do not result in "time bombs" which cause the response to be executed later unnecessarily.) Whenever all the fault responses originating from a given monitor are completed, that monitor and all monitors of the same or lower priority are initialized. An exception to this is the Command Loss monitor, which is only initialized when a command is successfully received. (The Command Loss persistence check can be extremely long, so resetting it could result. in a very long time to detect an actual command loss.) All monitors except the Command Loss monitor are also initialized, following a suspension of SFP while the CDS corrects a peripheral fault. However, following a boot of the CDS processor, all monitors need to be initialized. Initialization of responses consists of

resetting internal history parameters, e.g. counters which indicate how many previous response executions have occurred. All responses are initialized following a boot of the CDS processor, or following suspension of SF'P while the CDS corrects a peripheral. fault. A response is also initialized if it, or a response it requested, is cancelled due to a priority conflict-.

7. RULES AND PRACTICES OF THE DESIGN

Each monitor and each response has the capability to be enabled or disabled in either of two ways. There is a "S/C" enable/disable flag which can be changed by SFP algorithms, and there is a "MOS" flag which can be changed by "mission operations system" sequences, i.e. the mission activity sequences stored onboard. Either flag can be altered by direct ground command. in practice, the enable/disable states of responses and/or monitors are made appropriate to the current mission activity by the stored mission sequence, using the MOS flag. Responses or monitors which become inappropriate because of fault protection actions are. disabled using the S/C flag. In this way, the stored sequence cannot override an action taken in response to a fault. (Given that all responses have enable/disable flags, monitors only need to have this capability where more than one monitor can request the same response. However, in order to carry a consistent design, every monitor has been implemented with an enable/disable capability.)

By design, the enable/disable state of a monitor only affects its ability to request a response. Therefore, monitors will detect faults at all times, which is helpful in alerting the ground to unexpected conditions even though the fault. response may be deemed inappropriate at the time.

Monitors use up to three "filters" to eliminate bad data. Two of them, reasonableness and persistence checks, were mentioned earlier. The reasonableness check performed by a monitor is applied to analog data, such as temperatures, voltages, and currents . Usually this means declaring the data bad if it is at or near zero or full scale. This filter is not used if zero and full scale are va] id readings. Noisy data is eliminated by the persistence check, which requires a certain number of consecutive samples to indicate a failed condition before the response is requested. Any good sample, i.e. within acceptable limits, will reset the persistence counter. This means that a true failure could be masked by noise, and so the required persist-ence period is selected to be small enough to keep this possibility small . Persistence filter values are also picked to allow for reset recovery or warm-up periods and to ensure that a fault which produces two different symptoms is responded to correctly.

A third "filter", which makes use of redundant measurements, is used where the effect of a false detect ion is deemed severe. An example of a severe response is one which actuates pyrotechnic devices. Majority voting of measurements is preferred in these cases . In cases where only two measurement-s are available, and the

possibility of no response to a real failure is tolerable, both must indicate a failure before the response is requested.

All of the foregoing use of filtering notwithstanding, one of the primary considerations applied to response design addresses the possibility of a false alarm. Thus, it is requi reel that responses be designed such that an inadvertent execution of a fault response not be hazardous to the mission. Another rule, which can minimize the S/C impacts of a real fault. as well, is that if a fault response consists of multiple sequential. actions, the actions are executed in t-he order of increasing severity, even if this means that the most probable cause is not addressed initially.

Another requirement on response design is that it must be tolerant of being interrupted at any point without a harmful effect on the s/c. This accommodates t-he possibility t-hat. the CDS could suspend SFP if a CDS internal fault were to occur. The design is such that no uncompleted string of actions within a response places the S/C at risk if SFP is cancelled, then restarted from the beginning after CDS corrects its fault.

one final design rule is that modifiable parameters must be established where changes are reasonably expected to occur due to mission phase, inflight fault history, or actual performance of S/C hardware. These parameters are to be modifiable without. requiring a software reload or patching. Modification via ground commands is usually sufficient, except for a subset of parameters which must be changeable by the stored-mission sequence. This is necessary when it is impossible or unreliable for the ground to make a time critical change by realtime command.

8. PRIORITIZATION OF MULTIPLE CONCURRENT FAULTS/SYMPTOMS

The FP Manager implements the priority driven, one-at-time, architecture which handles cases where more than one monitor requests a response in one management: cycle, or where a monitor requests a response while a previously requested response is still active. This is implemented in the fol lowing manner. Each monitor is assigned to one of four possible priority levels. Responses which are in process, including those which have been requested by a "parent" response, retain the priority level of the originating monitor. During each one second management cycle, the Manager evaluates the monitor priority of any new response requests and any responses which are currently active. (For new response requests coming from monitors of equal priority, the monitor selected for evaluation is the one which is executed first by the software.) If there are no currently active responses, the response requested by the highest priority originating monitor (HPOM) is activated. If t-here is a currently active response, and the new HPOM i s of higher priority than the HPOM from the previous cycle, currently active response(s) are cancelled, and the response requested by the new HPOM is activated. If the priority level of the new HPOM is the same as that of the previous cycle's HPOM, the Manager takes no action, and the currently active response cent inues. If a

currently active response is cancelled, the Manager initializes it and enables its originating monitor via its S/C flag. This ensures that the first fault can be redetected, even if the cancelled response had disabled its originating monitor.

A number of design practices evolved from the fact that responses can be cancelled due to a priority conflict. As mentioned earlier, there was already the need to design responses so that they could be halted anywhere, due to the possibility of a CDS fault. When cancelled due to a priority conflict, however, the halt- takes place only at specified pause points in the response. Therefore, responses were designed with this in mind, e.g. activities were grouped to accomplish as much as possible, and in whole functional blocks, before reaching a pause point. Also, responses were designed to ensure that the original fault indication would still be present after completion of an interrupting response, if it was still necessary to accompl ı sh uncompleted acti ons of the cancelled response. To ensure that the effects of a cancellation would not be harmful, each response was subjected to a "redetect analysis" . In this analysis, a cancellation was assumed at. every pause point by each higher priority fault. The goal of the analysis was to show that the original fault would be redetected after completion of the higher priority response, or, if not, that the remaining actions would be accomplished by the interrupt-i-rig response or their omissi.on would be benign. The task of showing that uncompleted response actions were benign was made easier by the fact that, with only one exception, all responses capable of interrupting a lower priority response caused the safing response to be executed. The safing response i s designed to put the S/C into a safe state, regardless of its previous state.

The assignment of priority levels to SFP monitors was a process which started by making a few obvious assignments, and then used an iterative approach until a workable structure was attained. The highest priority was given to the most time critical faults, as would be expected. At the other end of the priority ladder, the lowest priority was assigned to the response with t-he longest. duration. The command loss response can run for several days, so it was important to not let its length preclude other responses from running. 'l'he sorting out of other monitors' priority assignments was based on considering a variety of factors: relative criticality, functional dependencies of t-he monitored hardware, and the amount. of similarity between different. responses, for example. Assigning non-unique priority levels involved difficult compromises, but many of the same considerations would have applied in assigning unique priority levels.

9. SPACECRAFT SAFING AND COMMAND LOSS DESIGNS

The following will describe design details of two Cassini SFP elements which are likely to exist in any S/C fault. protection design.

Spacecraft Safing

The first of these, S/C Safing, has the general purpose of placing the S/C in a safe, commendable state following a system level anomaly. in particular, the Safing response design has to assume the S/C can be in any state due to an interruption of t-he onboard stored mission sequence, as well as the unexpected state changes caused by the anomaly itself. One of the objectives of the Safing response is to satisfy the requirement to maintain a safe, commendable state following an anomaly for at least two weeks. The Safing response does this by configuring S/C! loads to a low power state, stopping any propulsive or attitude maneuvers, initiating the establishment of a robust attitude control hardware configuration, pointing the S/C to an attitude which both provides for solar thermal shading and allows earth communications, configuring telecommunications hardware to provide safe and robust link margins, and configuring electrical heaters to keep equipment within operational limits. The fact that stored sequences are halted by the FP Manager in al.] except critical phases of the mission ensures that the state established by Safing will. not_ be changed, except by other fault protection actions.

Safing does not by itself place the S/C in the minimum power state. Unless there is a power shortfall, Safing only places the S/C in a low power state. If the i.niti.sting event is a power anomaly which causes the total S/C load to exceed the available power, hardware logic in the power PPS sheds all loads which are not immediately essential . For example, heater loads are shed, whereas the CDS and AACS computers are not shed, although their peripheral hardware units are. In support of this hardware response, Safing, along with CDS and AACS subsystem fault protection, will power those loads which are essential for long term S/C health and commandabi lity.

The initial actions of Safing, after the FP Manager has stopped onboard stored sequences, are to stop propulsive and attitude maneuvers, then turn off al 1 science subsystems plus those heaters and engineering units which are non-essential . This immediately increases the S/C power margin to cover later Safing power increases, as well as t-hose caused by other SFP responses. 'l'here is a S/C operating margin which temporarily covers any power increases which occur prior to Safing. All. increases caused by subsystem fault protection are required to be within this margin.

At this point, the AACS subsystem is commanded to go to its safe state and to orient. the S/C to a thermally safe and commendable atti tude . (If required, the CDS autonomously reconfigures itself to its safe state prior to executing SFP, so the Safing response does not need to command this action.) When this is done, the Safing response turns power on to the prime telecommunications units, powers off their backup units, and selects antennas, modulation parameters, and data rates and modes to support uplink and downlink operations. Backup telecommuni cations units are powered off, primarily to prevent. interference with prime units, as well as to conserve power.

'I'he final Safing actions are to turn on essential engineering and science heaters. These actions reduce the S/C! power margin and are, therefore, placed at the end of Safing so that all temporary power increases due to other Safing actions will have completed. 'l'he heater powering actions are bypassed in the first few hours after launch, when full. power is not yet available from the S/C power source. The onboard stored sequence sets a flag which enables this branch of Safing at a predetermined time.

Command Loss

The second element of common interest to other S/C fault protection designs is Command loss fault protection. A small amount of protection is inherent in other responses, e.g. loss of power to the S/C receiver will cause a loss of exciter RF output and result in SFP commanding a receiver swap. The Command loss monitor and response, however, provide a comprehensive and uni fied approach to re-establishing command capabil ity following a loss caused by any S/C failure (or ground error) . The Command Loss monitor consists of a timer which counts down from a programmable value until it reaches zero or is reset. Upon reaching zero, the Command Loss response is requested. The receipt of a valid uplink command by the CDS will reset the timer to its original value and restart the countdown. Thu S , there is an end-to-end check on command functionality.

The Command Loss response contains actions which address all S/C related causes of non-commendability, inducting antenna failures, incorrect uplink antenna selection, incorrect. uplink rate selection, receiver/command detector failure, receiver interference from S/C RF sources, CDS hardware failure, CDS software anomaly, and incorrect S/C attitude. AACS health is assumed, since it is ensured by MCS subsystem fault. protection and by the MCS Heartbeat Loss monitor.

There is a basic strategy difference between the Command Loss response and all. other SFP responses, which derives from the seriousness of losing commendability. Whereas all other responses ensure recovery in the event of a single failure, the Command Loss response is designed to handle two recoverable failures. Thus , recovery is possible in the presence of a latent. recoverable fault, e.g. a previously undetected, but correctable, fault in a backup unit . The response accomplishes this by cycling through a sufficient. number (not all) of the hardware combinations. Thi s approach also handles non-viable combinations, where two otherwise healthy units do not work properly together.

The execution of the Command Loss response begins with the non-severe steps of resetting the command detector, selecting the auxi li ary osci 1 later in place of the radio science oscillator for the downlink frequency reference, and executing the Safing response. Because power cycling of telecommunications equipment carries a certain amount. of mission risk, no hardware swaps are

commanded until a ground response interval. elapses after these initial actions. Science instruments are turned off as a part of Safing, and this eliminates potential. RF interference sources in the science complement. In the next phase of the response, while using the safing antenna appropriate for the current mission geometry, seven new combinations of the receiver/cc)mmand detector, RF power amplifier, and the processor which cent.rc)ls telecom hardware modes are sequentially tried, allowing a five hour ground response interval between successive combination changes. At. the end of this phase, the original combination is re-established.

If the command loss condition is not corrected by this point-, the S/C! is oriented to the sun and put. into a slow roll about the sunline . This is designed to ensure a periodic telecom link, without relying on the star identification process in AACS . After seven hours (assures three revolutions), rolling is halted, and the alternate safing antenna is selected and pointed at. the earth for a thermally tolerable amount of time. This potentially risky orientation is required if the current safing antenna has failed, which is a small but non-zero possibility.

No CDS failure possibilities have been addressed up to this point, but before the response does so, it cycles through the hardware combinations one more time, this time varying the order so that when the S/C roll and alternate antenna actions are tried, it will be with the redundant hardware units.

Finally, the current prime CDS is commanded to reset, which will cause the backup to become prime. After its Command Loss timer expires (t-he two CDS'S timers may not have been in synch), the new prime CDS will re-initiate the Command Loss response from the start . If the original backup had been off at the first start of the Command Loss response, it will receive a software load from the onboard solid state recorder (SSR) when i-t is powered on. The original. prime CDS receives a load from the SSR when it. is reset, so if the command loss is due to a problem inadvertently programmed into the original memory load(s) , it. will be corrected either after one or two passes of the response.

10. SUMMARY

The Cassini system fault. protection design combines a standard monitor and response structure with a newly developed priority manager which precludes response conflicts. This design is the result of not only technical considerate ions, but. also concern over testability at the system and subsystem levels. Other features of the architecture have also been driven by both technical and practical considerations. Despite the design constraints which derive from these high level aspects of the architecture, all desired functions and features of the individual monitors and responses have been implemented successfully. Overall, the system fault protection design provides practicality, along with functionality and flexibility which are commensurate with the complexity and length of the Cassini mission.

This work was performed by the Jet Propulsion laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Biographical Sketch

John Slonski is a Senior Member of the [1'ethnical Staff in the Systems Division of the Jet Propulsion Laboratory. Previous work has been concerned with the spacecraft system design and development, plus flight operations, on the Magellan, Viking, and Mariner 1969 Projects. Currently he is t-he Deputy Spacecraft System Engineer for the Cassini Project. He received a bachelor of physics degree from Cal tech and a masters degree in electrical engineering from Stanford.

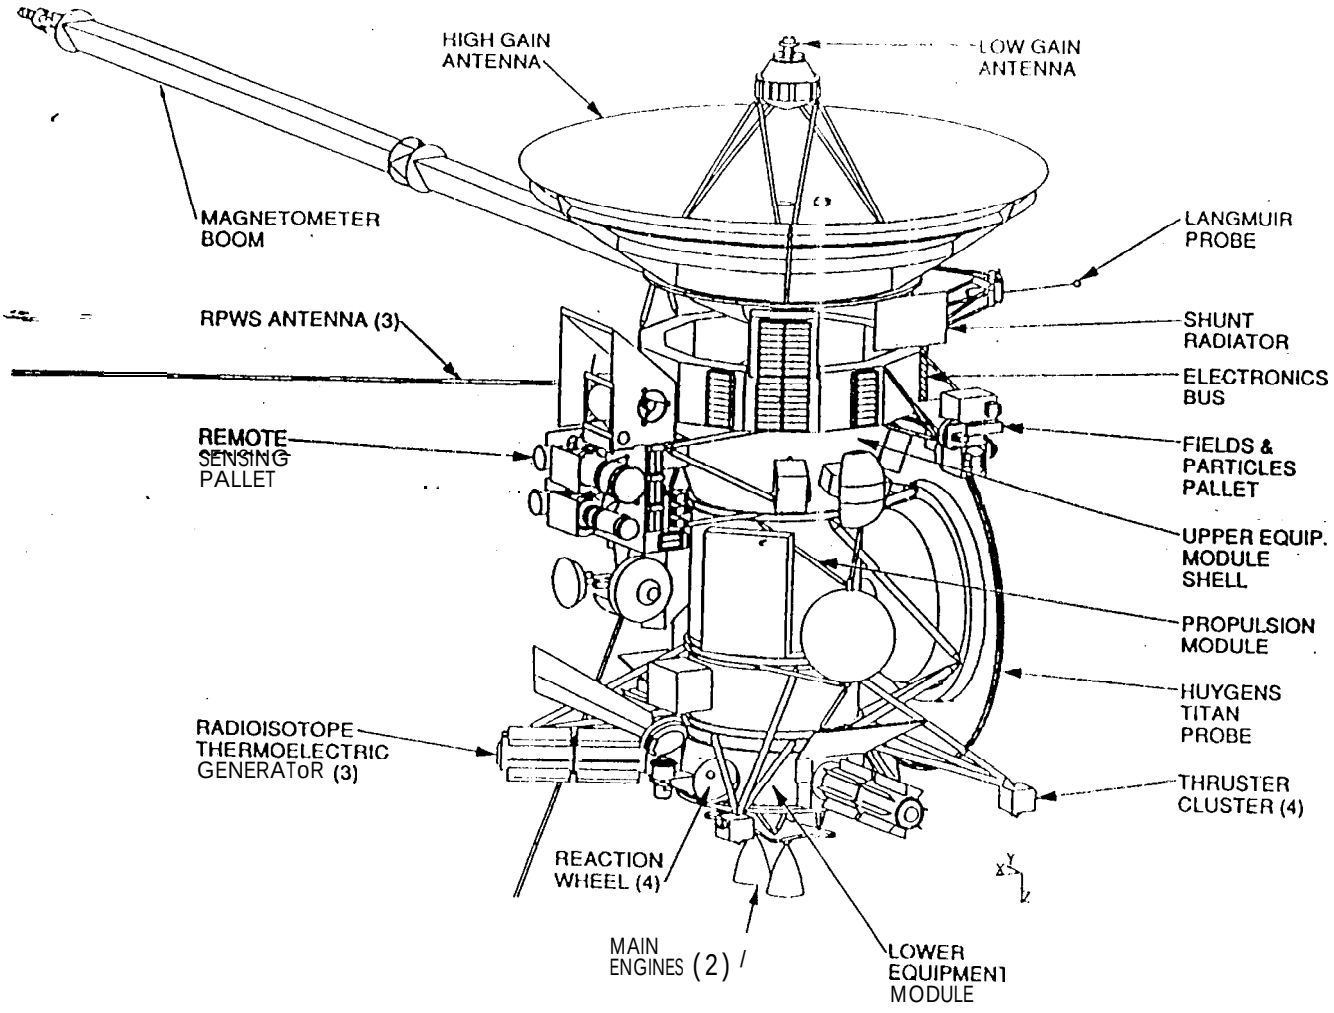

Figure 1. Cassini Spacecraft

Table 1

Single Point Failure

SPF Exemption #	Failure
1	Loss of a Radioisotope Thermoelectric Generator (RTG)
2	Loss of High Gain Antenna (HGA), or either Low Gain Antenna (LGA 1 or LGA 2) inside 1.5 AU.
	Leakage or bursting of a propulsion module tank (pressurant tank, main engine oxidizer tank, main engine fuel tank, thruster hydrazine tank)
4	External leakage or bursting of propulsion module fluid or pressurant lines and fittings, of components in the lines, and of pressure transducers. (Leakage past a closed thruster, engine or fill valve is not exempted).
S	Structure (Spacecraft adapter, orbiter, or Probe truss)
6	Spacecraft separation band (retention / release)
7	Thermal blankets, surfaces, and shields (spacecraft and probe)
8	Spacecraft cabling short
9	Selected command and data errors (*)
10	Main engine combustion chamber (catastrophic explosion)
11	Passive radio frequency equipment (3 db hybrid)
12	Micrometeoroid shielding (inherent or specific)
13	Power interruption greater than 37 msec
14-18	Probe adapter structures, Probe structure, spin-up and release mechanisms (exemption not applicable to premature release), heat shield, parachute system

)

)link commands:

- Untimely destruction of flight software or sequence memory through incorrect addressing or misuse of uplink commands.
- Untimely commands leading to an inappropriate subsystem state.

ı-board Erroneous Command Generation:

- Stuck bit in the bus controller memory address register leading to creation of a set of mission catastrophic bus transactions.
- Stuck bit in hardware command decoder message start detection circuitry leading to improper decoding and issuance of mission catastrophic commands

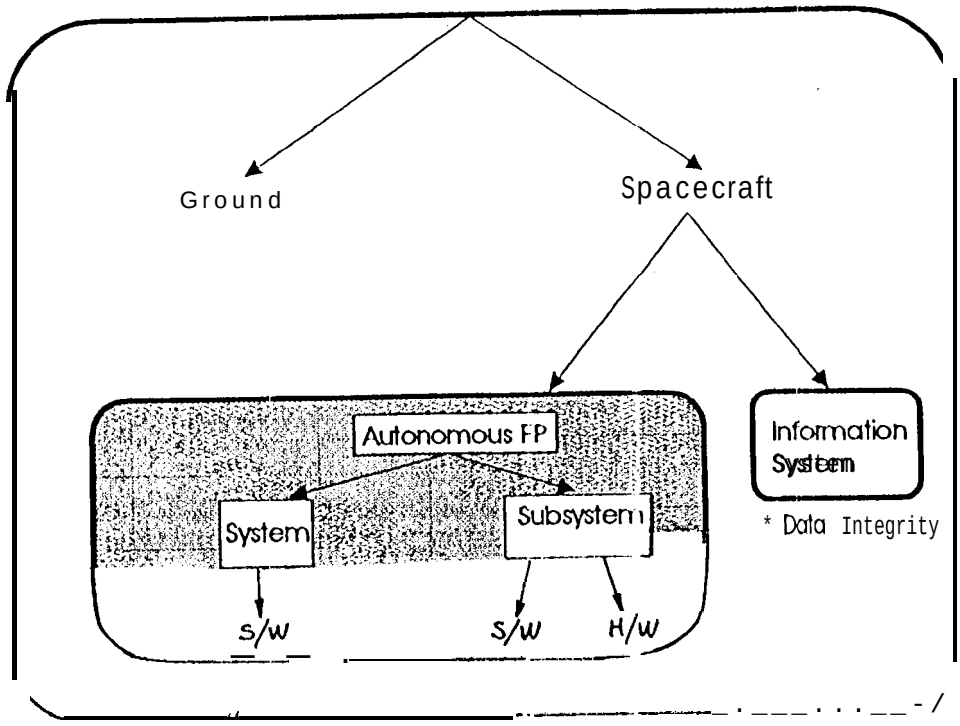

Figure 1. Fault Protection Responsibility Allocations

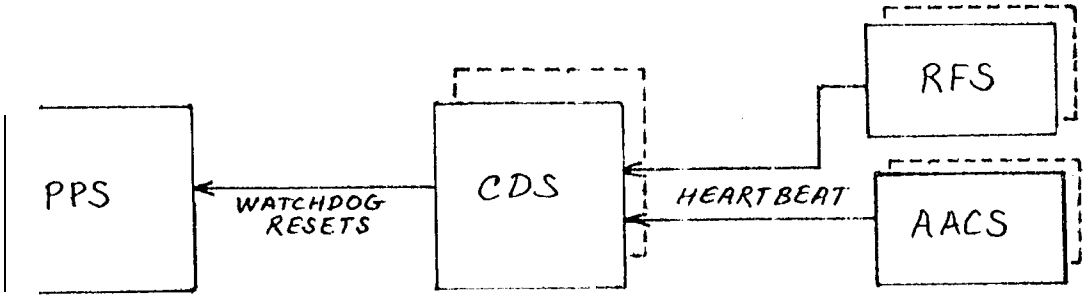

Figure 2. System of Processor Watchdogs

# MONITORS	# RESPONSES	PURPOSE
1. XCTR RF Loss 2. TWTA RF LOSS	1. XCTR RF LOSS 2. TWTA RF Loss	Recover from a loss of downlink
3. Command loss	5. Command Loss-J	Recover from a loss of uplink
4. RFS Heartbeat Loss	7. RFS Heartbeat Loss	Recover from a failed RFS TCU
5, Undervoltage	8. Undervoltage	Recover from a loss of power
6. Prop Tank Overpressure	9. Prop Tank Overpressure	Recover from a prop tank overpressure
7. Emergency O/T	10. Emergency O/T.	Recover from adverse thermal Increases
8. AACS Hearbeat Loss	11. AACS Hearbeat Loss	Recover from a failed AACS AFC
9. Alert Messages	12. AACS Safing Request 13. CDS Loss 14. Engine B Pyro	Support the CDS/AACS system FP Interface
	15. CDS S/C Safing	General S/C Safe

Figure 3. SFP Monitors and Responses